

Growing points in ethology

EDITED BY

P. P. G. BATESON & R. A. HINDE



1

Hierarchical organisation: a candidate principle for ethology

RICHARD DAWKINS

THE NEED FOR GENERAL PRINCIPLES: SOFTWARE EXPLANATION

If we look far into the future of our science, what will it mean to say we 'understand' the mechanism of behaviour? The obvious answer is what may be called the neurophysiologist's nirvana: the complete wiring diagram of the nervous system of a species, every synapse labelled as excitatory or inhibitory; presumably also a graph, for each axon, of nerve impulses as a function of time during the course of each behaviour pattern. This ideal is the logical end point of much contemporary neuroanatomical and neurophysiological endeavour, and because we are still in the early stages, the ultimate conclusion does not worry us. But it would not constitute understanding of how behaviour works in any real sense at all. No man could hold such a mass of detail in his head. Real understanding will only come from distillation of general principles at a higher level, to parallel for example the great principles of genetics – particulate inheritance, continuity of germ-line and non-inheritance of acquired characteristics, dominance, linkage, mutation, and so on.

Of course neurophysiology has been discovering principles for a long time, the all-or-none nerve impulse, temporal and spatial summation and other synaptic properties, γ -efferent servo-control and so on. But it seems possible that at higher levels some important principles may be anticipated from behavioural evidence alone. The major principles of genetics were all inferred from external evidence long before the internal molecular structure of the gene was even seriously thought about. Three computers with the same programming instruction set are in an important sense isomorphic in principle, even though their wiring diagrams may be utterly different, one employing valves, another transistors and the third integrated circuits; how all three work is best explained without reference to particular hardware at all (Simon, 1973). If a computer is doing something clever and life-like, say playing chess, and we ask how it is doing it, we do not want to hear about transistors, we simply

accept them. The useful answer to the question is purely in terms of software; indeed the programme is likely to be written in such a way that it could easily be run with completely different hardware.

We need *software explanations* of behaviour. I do not mean that animals necessarily work like computers. They may be very different. But just as the lowest level of explanation is not always the most appropriate for a computer, no more is it for an animal. Animals and computers are both so complex that something on the level of software explanation must be appropriate for both of them.

Ethology has not lately produced many general explanatory principles. Its energy models of motivation were bold and aesthetically satisfying, but the predictions they made were too simple, and they were vulnerable to the first attacks of powerful critical intellect (Hinde, 1956, 1960). In the aftermath of their destruction grand general principles understandably became unfashionable. Some ethologists switched to other problems, such as the ecological roles of behaviour, and I have even heard the word 'ethologist' used to label a man who was *not* interested in mechanisms of behaviour!

At Madingley the broad balance which is the best feature of ethology (Tinbergen, 1963) has never wavered. It is characteristic that the Sub-Department of Animal Behaviour should celebrate its quarter century with an imaginative look forward rather than by dwelling on the past. I have tried to show that if we look forward far enough into the future, we are driven to seek general principles rather than detailed minutiae. We may as well start thinking now about likely *candidate principles*, and before looking at wholly new ideas it is worth dusting off some of the old ones. Hierarchical organisation is considered central to the whole of biology by the founder of the Sub-Department (Thorpe, 1974). Its ethological manifestation (Baerends, 1941; Tinbergen, 1950, 1951; Hinde, 1953; Kortlandt, 1955) came to grief in the general, deserved destruction of simplistic energy models, but its guilt was by association only. It really has nothing to do with energy models, but is a much more powerful principle in its own right. It is a particular pleasure to a pupil of Niko Tinbergen to try to point this out.

This paper will not be a review of the literature, nor will it use hard evidence to convince anyone. It will be an attempt to arouse the imagination of those more accomplished in research than I am, to persuade them to look again at the idea of hierarchical organisation, and use it in the future. But hierarchy is only one example of a principle of software explanation, and if it is eventually found wanting, other possibilities may be explored in a similar way.

HIERARCHICAL ORGANISATION

DEFINITIONS AND CLASSIFICATION

ἱεραρχεῖν means to be supreme in sacred things (Liddell & Scott, 1883). The idea of supremacy or superiority is the basis of the following definitions, which are indented to distinguish them from plain English. They will be best understood with reference to Fig. 1. First the axioms:

There exist elements: A, B, C etc., and a relation:
'is boss of' (inverse: 'is bossed by').

The elements can be represented by blobs, and the relations between them by arrows (Fig. 1). If A is boss of B then an arrow is drawn from A to B. We now define a more general relation:

A is *superior* to B when either
A is boss of B or
A is boss of an element which is superior to B.

Thus an element can be superior to another element without being boss of it. Colloquially boss might mean 'immediate superior'. Using this preliminary definition we can now define a hierarchy. (I prefer to define it as a set of




Fig. 1. (a) Linear hierarchy ('peck order'). (b) Non-overlapping branching hierarchy. (c) Overlapping hierarchy. (d) Not a hierarchy. (e) Shallow hierarchy. (f) Hierarchy of loops. (g) Network.

elements with a relation, rather than to follow Woodger (1937) in defining it as a class of relations. My usage is closer to that of the Church, in which the Hierarchy is the set of bishops.)

A *hierarchy* is a set which satisfies:

- (i) There is no element in the set which is superior to itself, and
- (ii) There is one element in the set, to be called the *hierarchy*, which is superior to all the other elements in the set.

The first requirement (if strictly applied – see below) is equivalent to saying that the relation 'is superior to' must be transitive throughout a true hierarchy, i.e. there are no circular relationships as in Fig. 1d. The second requirement ensures that every hierarchy has a single overall superior or root. There is no reason why the hierarchy should not be bossed by another element, but that element is by definition outside the hierarchy under discussion. Thus a hierarchy may be included in a larger hierarchy. Hierarchies may be classified into types as follows.

A hierarchy is *branching* if it includes at least one element which is boss of more than one element.

A hierarchy is *linear* if it is not branching.

A (necessarily branching) hierarchy is *overlapping* if it includes at least one element which is bossed by more than one element.

The model of a linear hierarchy (Fig. 1a) has been applied to social 'peck orders', and it has also been used, not always under the same name, for relationships within animals between different behavioural subsystems (Deutsch, 1960; Dawkins, 1969b; Davis, Mpitsos & Pinneo, 1974; McFarland & Sibly, 1975). It is a worthwhile candidate principle, but in this paper I am concerned mainly with branching hierarchies, both overlapping and non-overlapping. Indeed the very word hierarchy carries connotations for many people of a branching tree. Equivalent representations include nested brackets and overlapping areas.

The branching hierarchy idea has found its way into biology many times (Weiss, 1971; Pattee, 1973), and it is discussed by systems theorists (Mesarovic, Macko & Takahara, 1970). For reasons which I do not understand – perhaps it has something to do with the ecclesiastical meaning of the original Greek – it is a favourite of mystics (Koestler, 1967), or at least of those who are fond of using the word 'reductionism' (Koestler & Smythies, 1969). It is the basis of both Linnean and post-Darwinian taxonomy, although many people persist in seeing the animal kingdom as a linear hierarchy (Hodos & Campbell, 1969). It may have important applications in ontogeny at

various levels (Manning, this volume; Simpson, this volume). As Kortlandt (1955) has shown in a brilliantly erudite, if somewhat outspoken review, branching hierarchy models have a long history in human and animal psychology, pre-dating the more familiar ethological ideas.

The nature of the relation 'is boss of' has been left unspecified so far. Interpret it as 'feeds motivational impulses to' and you have Tinbergen's (1950) hierarchical model. Change it to 'has a causal influence on', and you arrive at Hinde's (1953) less contentious formulation:

... of the causal factors controlling each activity, some are specific to it and others are also capable of influencing other activities. The causal factors are thus arranged in a hierarchical manner, those at the bottom influencing only one or a few activities whereas those at the top influence many. Ultimately, of course, there are factors which influence the occurrence of all possible behaviours – these are the conditions under which existence is possible... This hierarchy of causal factors gives us one way of classifying the various activities possessed by an animal into instincts and sub-instincts. All the activities belonging to a major instinct have certain causal factors which they share with other sub-instincts and some which are specific to them.

To list some other examples of hierarchies, 'is boss of' might mean 'controls' (as in a power station), 'is an ancestor of' (as in a family tree), 'is attached to' (as in an oak tree) or 'gives orders to' (as in the army). As we shall see, several interesting behavioural models are based on its meaning 'sets up a target or goal for'. In a computer programme it might mean 'calls up as a subroutine (procedure)'. A related meaning will be shown to give rise to another interesting class of models, those comparing behaviour organisation to human grammar.

What would *not* constitute a hierarchy? Within the context of sets of elements with a single relation 'is boss of', it follows from our definition that there are only two classes of non-hierarchical system, those which violate conditions (i) and (ii) on page 10. These are non-transitive systems and multi-hierarchy systems respectively. To these might be added systems like Fig. 1e, in which there is a hierarchy ruling over many subordinates, but all the subordinates are of equal status to each other. These 'shallow hierarchies' do not strictly violate the definition, but they are not hierarchical in an interesting sense, because their 'height' (Harding, 1971) is small compared with their 'span' (Simon, 1962).

Fig. 1g taken as a whole violates the single hierarchy clause. Some people are reluctant to give the name hierarchy to overlapping hierarchies in general, because they feel that a logical extension of an overlapping hierarchy is a network like Fig. 1g, with no evident tree-like properties. However, although it is true that the set of all elements in that figure is not a hierarchy, it is possible to regard it as a set of hierarchies. Thus if all arrows are regarded

as pointing downwards, each element in the top row becomes a hierarchy superior to nine lower ones. One tree is emphasised in the figure for illustration. It is in this sense that a *Wirkungsgefüge* (von Holst & von Saint Paul, 1963) is hierarchical. Hinde (1952) pointed out long ago that even Tinbergen's model must really be overlapping, at least at the lower levels. The same is clearly implied in Weiss (1950) and in Lorenz's (1937) concept of 'tool-instinct'.

A network such as Fig. 1g functions as a collection of true hierarchies in the following sense. Although two hierarchies share many subordinates, each one has a unique set of them. Thus two cells in the visual cortex of a cat may have overlapping fields, both retinal and supra-retinal, but they will not overlap totally. Hence each one will be maximally stimulated by a unique pattern, though they share subpatterns, features, or at least spots of light. A similar point has been made for the motor side, for example with reference to invertebrate 'command interneurons' (Elsner, 1973; Kennedy, 1974).

Should systems with feedback be regarded as non-hierarchical because they violate the transitivity rule? Not necessarily. It depends how we interpret the word 'element' in the definition. In Fig. 1f, if the black blobs are the elements, the set is clearly not a strict hierarchy. But if we put each feedback loop into a box, and call each box an element, we then have a transitive hierarchy of feedback loops. Thus the TOTEs of Miller, Galanter & Pribram (1960) form true hierarchies. Ito (1974) postulates 'pyramid-like' hierarchies of control in the nervous system, with feedback units at lower levels, and feedforward and other kinds of control at higher levels, with an overall controller at the peak. Szentágothai & Arbib (1974) provide a stimulating discussion of feedback and feedforward in hierarchical systems.

On the other hand if feedbacks are not locally confined, but extend to more and more distant superiors it becomes less and less easy to regard the system as a hierarchy. There is no need to specify a definite cut-off point for the definition. This would be no more sensible than worrying about whether two animals belong in different genera or merely in different species.

However, simple rings such as Fig. 1d should definitely be ruled out. An example is a 'round-robin' computer time-sharing system. The different tasks are allocated a fixed share of the computer's time in strict rotation, like a chess master playing ten games 'simultaneously'. In less efficient round-robins, the timing is inflexible, even if particular tasks cannot make use of their time allocation for want of a busy peripheral device. R. H. McCleery (personal communication) suggests that *Arenicola* (Wells, 1966) may organise its behaviour like this. A more efficient method of time-sharing, a linear hierarchy by our definition, is 'priority-interrupt'. The top priority task retains control until it is held up, usually waiting for a slow peripheral device or human,

whereupon the second-ranking task is disinhibited. Number two retains control until either it is held up and disinhibits number three, or it is interrupted by number one. McFarland (1974) has developed the analogy with animal behaviour.

Pure priority-interrupt systems have the potential disadvantage that very low ranking programmes may never be run at all. This problem can be overcome by a compromise with a round-robin, or by shifting priorities as a function of time, much as an animal gives higher priority to eating as the interval since its last meal increases.

Shifting of boss-ships among a community of potential equals is the basis of common ethological motivational ideas (e.g. van Iersel & Bol, 1957), in which everything inhibits everything else; the 'centre' which is for the moment most potent in inhibiting rivals, takes control of the final common path. By analogy with the 'pandemonium' system of sensory pattern recognition (Selfridge & Neisser, 1963), this may be called the 'motor pandemonium' model: whoever is shouting loudest at the moment is boss. Physiologically speaking, fluctuating factors such as hormonal and nutritional state, levels of skin irritation and other sensory factors, change the priorities among the competitors. Changes in observed behaviour are all directly caused by these changes in boss-ships. This is not true of interestingly hierarchical models such as McFarland's time-sharing model (see above), in which 'is boss of' means 'is capable of interrupting', and my 'attention threshold model' (Dawkins, 1969b), in which 'is boss of' means 'is attended to before'. In both these models, boss-ships remain stable while behaviour changes. The motor pandemonium model is not hierarchical except in a trivial sense. It is shallow (see above), and its slight hierarchical properties are not used to explain anything.

Nelson (1973) espouses 'distributed control', at times apparently setting it up as anti-hierarchical. Unfortunately, in a paper so rich in creative ideas, he does not find the space to define it. He is impressed with the power of simple networks of very few neurone-like elements to generate complex output patterns, and to switch abruptly from one stable pattern to another in response to simple quantitative changes in input. He goes on to visualise more complex networks in which 'the organization of connection would still be in a way hierarchical but lacking a hierarchy. Control or dominance would shift from one to another part of the distributed perceptual-motor interface according to internal or external necessity. . . .'

McCulloch (1945), also in the context of distributed control in networks, used the term 'heterarchy', and the word has recently been adopted by workers in artificial intelligence (Winston, 1972). Unfortunately it too has not been defined, as Winston admits. He goes on to provide a complicated list

of five unconnected attributes of heterarchy, of which only the first two are easy to generalise from computer programming to biology. First, heterarchical programmes are 'goal-oriented'. We postpone discussion of goals until later. Secondly, 'executive control should be distributed throughout the system. In a heterarchical system, the modules interact not like master and slave, but more like a community of experts'.

Nelson's 'distributed control' and the 'heterarchy' of the artificial intelligence workers appear to have something in common. In both, as in the motor pandemonium model, boss-ships are capable of rapid reversal. Does this make them fundamentally anti-hierarchical? I prefer to think of changing boss-ships as an interesting complicating factor in systems which may or may not be hierarchical according to other criteria.

For example junior officers take orders from their superiors, but they also send information back in the reverse direction. There is flow of orders 'down' the pyramid, and flow of information 'up'. We could think of shifting control, depending on whether information or orders are being considered. But there is a more interesting sense in which the system remains a hierarchy, because information tends to converge towards an apex, and orders to diverge from the same apex. A captain can be identified as superior to a lieutenant whichever the direction of flow, because a captain interacts with several lieutenants, but a lieutenant only interacts with one captain.

Even a motor pandemonium model can be extended, so that each participant can be the head of a stable hierarchy. In a network such as Fig. 1g, control might shift from point to point, so that new hierarchies emerge. The hierarchies so brought into play may still have the attributes of a branching tree, with all that follows functionally (see below), even though they are temporary. Fig. 1d and 1e, however, are not interestingly hierarchical, whether they are regarded as temporary or permanent.

Hierarchical classification

Nelson (1973) makes an important distinction between hierarchies of embedment and hierarchies of connection. I think embedment is largely equivalent to classification, in which case the same distinction has been emphasised by others in the ethology literature (e.g. Kortlandt, 1955). The examples given above were hierarchies of connection. A hierarchy of classification is one in which 'is boss of' has some such meaning as 'contains' or 'includes'. Thus the general feature of a hierarchy of classification is that inferior elements actually make up the parts of higher elements. This is not so of hierarchies of connection. To paraphrase Nelson, the Curia is not a partitioning of the substance of the Pope. A soldier is a part of his platoon, but not of his platoon commander.

Hierarchical classification is a vital convenience of everyday life, a universal means of organising information for easy access. A large postal system would be unworkable without hierarchical addressing. Even numerical ('zip code') addressing is really hierarchical, and in the same sense as the alphabetical arrangement of a dictionary (Longuet-Higgins & Ortony, 1968). Non-hierarchical methods of information retrieval such as random search, and sequential systematic search, are much too slow.

It is of course important not to muddle up hierarchies of classification and of connection. However, the muddling that has undoubtedly occurred in the ethological literature is perhaps pardonable, because the two are frequently closely associated. Systems such as armies are hierarchically classifiable into functioning units, but these units are controlled by a hierarchical tree of command or connection. Similarly, biological taxonomy is hierarchical not only because an embedded structure is convenient, but also because of evolutionary connections. Each node in a taxonomic tree-diagram can be regarded as a taxon such as Mammalia which can be decomposed into subgroups, Carnivora, Rodentia, etc. It can also be regarded as a common ancestor.

In ethology it is clearly convenient to *classify* behaviour patterns hierarchically for the same reason as it is convenient to classify anything else in this way. But, at least until the concluding discussion, we shall be more concerned with the other question of whether there are hierarchies of connection between the subunits. This could be discussed from a neuroanatomical and neurophysiological point of view (Paillard, 1960; Szentágothai and Arbib, 1974; Eccles, 1975). In this paper I shall take two more indirect approaches. The first is a functional one – it might even be called a 'neuroeconomic' one – an attempt to show some of the economic pressures which might lead to the evolution of hierarchical systems. The arguments here are sufficiently general to apply to hardware or software. The second is a behavioural one, an examination of the temporal patterning of behaviour to see whether it appears to be governed by hierarchical pattern-generating (software) rules.

THE FUNCTIONAL SIGNIFICANCE OF HIERARCHICAL ORGANISATION

An interest in the functional significance or survival value of a biological feature is often regarded as confined to field-workers who see whole animals in their natural surroundings. This is obviously silly. Natural selection entitles us to expect with confidence that good, though not necessarily optimal, design principles will pervade the internal organisation of animals down to the smallest levels. What is less obvious is that it is good research strategy to think about design principles before, rather than after, attempting

to find out what is actually the case. Good physiologists know this, and are guided by it straight to the most fruitful hypotheses for physiological testing.

Simon (1962), in a stimulating short paper suggesting that hierarchy is the basis of the 'Architecture of complexity', illustrates one reason why hierarchical design is good design by a parable about two watchmakers called Tempus and Hora. (Koestler in his re-telling (1967) changed the names to Mekhos and Bios – why?) Tempus's watches were as good as Hora's, but he took about 4000 times as long to make each one. The reason lay in a fundamental difference of design. Both types of watch had about 1000 components. Hora put these together first into 100 sub-assemblies of 10 components each. These in their turn were assembled into 10 larger units. Finally, the 10 larger units were combined to complete the watch. Thus if anything went wrong during assembly, Hora had only to go back and re-assemble the current subunit, and he did not lose much time. Tempus on the other hand tried to put together all 1000 components in a single large assembly operation. If there was a mistake or interruption the whole thing fell to bits, and he had to go right back to the start. He therefore very rarely completed a watch.

Parables should not be flogged to death, so I will refrain from relating stories about two computer programmers, or two sticklebacks, called Tempus and Hora, and turn to a listing of what seem to me the three main advantages of hierarchical organisation. When I wrote this I was unaware of the parallel but different discussion of Szentágothai & Arbib (1974).

The evolutionary rate advantage

This is the one illustrated by the watchmakers. In more abstract terms, the evolution of thermodynamically improbable assemblies proceeds more rapidly if there is a succession of intermediate stable sub-assemblies. Since the argument can be applied to the manufacture of each sub-assembly, it follows that highly complex systems which exist in the world are likely to have a hierarchical architecture (Simon, 1962), and nervous systems are presumably not exceptional. The word 'advantage' is perhaps misleading here, as indeed it often is in evolutionary discussion. 'Evolutionary stability' (Maynard Smith & Price, 1973) is nearly always better.

The local administration advantage

Consider the problem of controlling an automatic vehicle surveying Mars. The question is how much of the total electronic and computing power to put on Mars and how much to leave on Earth. One extreme, which has economic

appeal because of the low rocket payload, is to leave almost all the decision-making circuitry on Earth – a general purpose computer could do the job – and equip the vehicle itself with little more than a two way radio set. That the economy was false would soon be apparent to anyone listening in to the radio messages. Every time the vehicle met a small local difficulty, a boulder say, it would relay the details to Earth, each bit of information taking four minutes to arrive. In a flash the giant computer would calculate the optimum tactic, but each bit of the returning instructions would take another four minutes to reach Mars, and the wretched robot would long since have ploughed into the boulder. Obviously detailed moment to moment radio-control from Earth is prohibited by the delays. Moreover, since much of the information necessary for control is all on Mars in the first place, it is a waste of the communication channel to refer it back to Earth where it is not used.

Clearly many detailed decisions based on local information are best taken locally, and this is a fundamental principle of far-flung organisations such as the late Roman and British Empires. On the other hand the main disadvantage of too much local responsibility is lack of coordination of different units towards a common purpose. The optimum balance between local responsibility and referring back to headquarters depends in a complex way on a number of factors, including the distance involved, measured in rate and cost of information transmission. Call this the 'information distance'.

Now suppose we need to control several vehicles on Mars. The information distance between each vehicle and Earth is the same. But the vehicles need to be coordinated together in a common plan, and the information distance between any two points on Mars is relatively small. Therefore the optimum balance will tend to shift towards setting up a local master computer on Mars which will take some decisions for all vehicles, and which will handle all communications with Earth. We have the beginnings of a branching tree, and the argument can be applied recursively to justify further, many-levelled branching.

The redundancy reduction advantage

A classic of the discipline which I am naming 'neuroeconomics', is Attneave's (1954) and (independently) Barlow's (1961*a, b*) analysis of principles underlying sensory systems. Most messages contain redundancy, that is they could be recoded more economically without loss of information. For example in most visual scenes there is a high correlation between the luminosity of neighbouring points. This means that if the intensity of the light falling on each retinal cell were simply mapped onto the visual cortex, the firing rate of any one central projection cell would be highly predictable from its

immediate neighbours. This is not only wasteful of channel capacity; it is also unhelpful to the animal, which has to make practical decisions, to have information simply *reproduced*, however accurately, on yet another projection screen in its nervous system. As Barlow points out, nervous systems in fact recode so as to remove redundancy, lateral inhibition in the retina being the mechanism in this case. This ensures that maximal firing rates occur in cells bearing much information, which here are cells whose fields lie along *edges* in the scene. Similar arguments can be made for other aspects of sensory systems, and I will apply them also to the motor system. First we must make the link with the idea of hierarchy.

Many of the sensory patterns which an animal has to recognise – food, mates, obstacles, etc. – have features in common. This is another form of redundancy. Thus a straight line is redundant not only in Barlow's sense that all points along it are predictable from the two ends. It is also redundant in that it is a feature common to many ordinary objects. Rather than have entirely independent circuitry to recognise each important object it is therefore economical for pattern-recognising units to share subcomponents which recognise subfeatures which their key stimuli have in common. Recursive application of this argument leads to a system of overlapping hierarchical pyramids. This functionally sensible design is familiar from the work of neurophysiologists on visual mechanisms themselves, and of computer programmers facing the analogous problem of machine recognition of visual patterns (Sutherland, 1969; Barlow, Narasimhan & Rosenfeld, 1972).

Attneave and Barlow began their arguments independently by the same numerical thought-experiment. The human retina contains about four million light-sensitive cells. If we make the simplifying assumption that at any instant each cell is either signalling presence of light or not, the number of possible states of the system is $2^{4000000}$ which is not a small number. If there were one central cell tuned to each possible state of the retina, the volume of the brain would have to be measured in cubic light years. It was in this context that Barlow and Attneave postulated the pressing need for redundancy reducing mechanisms in the visual system.

Can we make a similar argument for the motor system? There are fewer muscle fibres in the human body than retinal cells, but there are enough to make difficulties. If we assume that at any instant each one is in a state of either contraction or relaxation we can arrive at a similar combinatorial explosion if we try to calculate the total number of possible states of the muscular system. In fact the potentially enormous number of states is greatly reduced by redundancy in the final motor output – large populations of muscle fibres contract and relax in a highly correlated way, and it is obviously sensible that they should do so.

The most perfect correlation is between members of the same 'motor unit', fibres which are all controlled by the same motoneurone. The correlation between motor units within one muscle is not perfect – if it were, graded contraction would be impossible. Nevertheless there is great redundancy in the behaviour of motor units within a single muscle, which is thus to some extent a unit of action. Then there are correlations between different muscles, both positive and negative, simultaneous and time-lagged. We are of course mounting the ladder of Weiss's (1941) well-known six levels of nervous organisation, which came to form the lower rungs of Tinbergen's hierarchical model. This arrangement confers the same economic advantage on the animal as subroutines give the computer programmer. Whole patterns of low-level coordination, programmed only once, may be called into the service of different high-level tasks.

Here then we have another economic argument in favour of hierarchical organisation from the animal's point of view. A similar, logical rather than functional argument can be advanced for the a priori plausibility of the nervous system's being hierarchically organised. This was briefly mentioned by Craik (1943) but it is fully set out by Bullock (Bullock, 1961; Bullock & Horridge, 1965). He considers 'the problem of recognition in an analyzer made of neurons'. He is concerned only with those cases where a definite all-or-none behavioural act emerges from the animal. How large this category is is open to empirical test (Dawkins & Dawkins, 1973), but many ethologists seem to assume in practice that it includes the behaviour they are studying. Bullock's point is that in such cases there must be, somewhere in the nervous system, a single unit which makes the decision. This unit may be regarded as the point of convergence of information from many sources, including sense organs, or it can be regarded as the starting point of outwardly radiating efferent information. In other words it is the confluence of two hierarchies. Bullock is much too cautious to call his units single cells, though this is one possibility he considers, and Barlow (1972) has recently argued provocatively for a 'neuron doctrine for perceptual psychology'. Certain large invertebrate neurones (Dorsett, Willows & Hoyle, 1973; Kennedy, 1974) would fit the bill. Bullock also considers multi-neurone decision-making units, which may go some way towards allaying the scepticism of Nottebohm (1970).

Hierarchy then seems to fulfil a major requirement of a good general candidate principle for the organisation of behaviour; it makes functional and logical sense. At present the evidence that it actually is an important principle is not convincing: not enough work has been done on it, which is one reason for writing this paper. I shall now turn to behavioural models.

SIMPLE CLUSTERING IN TIME

Eibl-Eibesfeldt (1975), who has given the most sympathetic recent treatment of Tinbergen's classic model, cites as evidence for it the fact that behaviour patterns tend to occur clustered in time, the clusters constituting functionally related groups. Tinbergen's (1950, 1951) own words can be interpreted to mean the same thing, although they have also been interpreted (and criticised – see above) as purely taxonomic in intent. Tinbergen of course also made use of other evidence such as that from electrical stimulation of the brain (see Vowles, 1970, for a more recent usage of the same kind of evidence in the service of another hierarchical model). In this section we are concerned with the relevance of the grouping of behaviour patterns in temporal clusters.

Simon (1962) considers a similar point more generally, under his heading of 'near-decomposability', for him a property of hierarchical systems generally, although he does not deal with animal behaviour. A small digression is needed to explain this.

Let the elements in a system be listed as the column and row headings of a matrix and let the body of the matrix contain numbers representing strengths of interaction between them. The order in which the elements are arranged as row and column headings is at first arbitrary. We now rearrange them so as to maximise the tendency for high interactions to be grouped in square submatrices around the major diagonal. The matrix is said to be decomposable if it is possible to arrange it so that all interaction scores lie in these square submatrices.

Fig. 2a, taken from Simon (1962), is an example of a *nearly* decomposable matrix. The elements around it refer to cubicles in a house, and the numbers in the table represent rate of heat flow between cubicles. Here it is possible to arrange the table so that all high numbers lie in three square submatrices along the diagonal, and all low numbers lie outside these submatrices. A rationale for the tendency to decompose into three submatrices is the following hierarchical scheme (Fig. 2b). The cubicles A1, A2 and A3 are in one room, B1 and B2 in another, and C1, C2 and C3 in a third. Insulation between rooms is good, while insulation between cubicles is poor.

More generally, low-level elements in a hierarchy are bound to each other by strong bonds, and their dynamic interactions are of high frequency (Bastin, 1969; Simon, 1973). Weaker bonds and slower dynamics characterise interactions between higher-level elements, which are clusters of low-level elements. If the differences between interaction strengths at different levels of organisation are large, the near-decomposability will show up in a matrix such as Fig. 2a, and the system can be treated as hierarchical.

Ethologists also represent strengths of interactions in square tables of this

HIERARCHICAL ORGANISATION

(a)

	A1	A2	A3	B1	B2	C1	C2	C3
A1	—	100	—	2	—	—	—	—
A2	100	—	100	1	1	—	—	—
A3	—	100	—	—	2	—	—	—
B1	2	1	—	—	100	2	1	—
B2	—	1	2	100	—	—	1	2
C1	—	—	—	2	—	—	100	—
C2	—	—	—	1	—	100	—	100
C3	—	—	—	—	2	—	100	—

(b)

A1	B1	C1
A2		C2
A3	B2	C3

Fig. 2. Hypothetical nearly decomposable system. (a) Table of heat diffusion coefficients between cubicles, arranged so as to concentrate high entries in square submatrices along major diagonal. (b) Plan view of cubicles, to show their clustering in three rooms, A, B and C (from Simon, 1962).


type. Sometimes the figures in the table are temporal correlation coefficients between behaviour patterns, and sometimes they are probabilities or frequencies of transition. The ordering of rows and columns is usually arbitrary. However, Myrberg (1972) published transition matrices for behaviour patterns of the fish *Eupomacentrus partitus*, in which he rearranged the rows and columns so that the entries in the body of the table came to be clustered in square sub-tables, although he did not provide an explicitly hierarchical rationale. It is not obvious whether he achieved the optimal rearrangement, and he does not give in detail the algorithm he used. However, Fig. 32 of Myrberg's monograph is reminiscent of Simon's figure reproduced here, and there is a suggestion of the near-decomposability which Simon regards as a property of hierarchical systems. A more exacting test of the near-decomposability of Myrberg's table is given on p. 31.

This method of rearranging rows and columns of a behaviour transition-matrix is a form of cluster analysis, but rather a crude form. It treats behaviour patterns as arranged in a two-levelled hierarchy, with strong interactions between elements within a low-level cluster, and weak interactions

(a)

	Following									
	FR	TG	HD	FM	NO	MV	BM	AB	WG	BF
Preceding	FR	77	709	129	496	5				
	TG	82			2					
	HD	730	4		18					
	FM	151		1	13					
	NO	445	3	42	36	418	26	133	495	414
	MV	7				412				1
	BM						13			76
	AB						93		3	195
	WG						223	1	1	46
	BF						741	62	92	156

(b)



(c)




Fig. 3. Data on grooming behaviour in blowflies, *Calliphora erythrocephala* (from Dawkins & Dawkins, 1976). (a) Table of frequencies of transition between acts, arranged so as to concentrate high entries in two submatrices. (b) Single Linkage Cluster Analysis (Ross, 1969). Index of 'distance' between pairs of acts was $(1-r)$, where r is product moment correlation between time spent doing the two activities in successive five second periods. Two acts are clustered close together if they tend to occur in the same 5 second period. Two existing clusters are united, if any one

HIERARCHICAL ORGANISATION

between elements of different clusters. Other kinds of cluster analysis are designed to group elements in multi-levelled hierarchical classifications (Everitt, 1974). Fig. 3b shows the application of one such, the Single Linkage method (Gower & Ross, 1969; Ross, 1969) to unpublished data on grooming in blowflies *Calliphora erythrocephala* (Dawkins & Dawkins, 1976). For comparison a table of transition frequencies is given with rows and columns rearranged in the same way as Myrberg's tables (Fig. 3a). Ignore Fig. 3c until later.

However, cluster analysis is only a technique of classification. As explained above, we want to know whether the underlying control is hierarchically organised. General cluster analysis techniques will tell us which of many alternative hierarchical clusterings best fit our data; they will not tell us whether our data really 'want' to be classified hierarchically at all, in the sense that the underlying mechanisms are hierarchically organised.

Similarly, we cannot necessarily regard the fact that behaviour patterns tend to be clustered in time, as evidence in itself for underlying hierarchical organisation. Behaviour patterns doubtless do form clusters in that each act is likely to be followed by another member of the same cluster, i.e. between-cluster transitions are rare. But, unless further conditions are met (see below), there is nothing here that could not be predicted by an ordinary one-levelled state-transition Markov model (references in Slater, 1973). Nothing is gained by speaking of clusterings of behaviour patterns. 'Rare transitions' and 'common transitions' are more parsimonious ways of expressing 'between-cluster' and 'within-cluster' transitions. We can indeed use such a Markov process as a kind of null-model, an example of a non-hierarchical model.

A behavioural model which allows us a truer test of near-decomposability is best introduced in terms of the concept of 'decision'. This must now be explained, since the word is used in a slightly unusual way.

of one cluster is close to any one of the other. (c) Mutual Replaceability Cluster Analysis (method explained later in text). Two acts are clustered together if they are mutually replaceable in behaviour sequences. Entries from existing clusters are added together for consideration for future clustering. Figs. 3(b) and 3(c) use the data in different ways, and are not simply mutually translatable. Key to behaviour patterns: 1. FR Rubbing front legs together; 2. TG Grooming proboscis with front legs; 3. HD Grooming head with front legs; 4. FM Grooming one middle leg with front legs; 5. BM Grooming one middle leg with rear legs; 6. BF Rubbing rear legs together; 7. AB Grooming abdomen with rear legs; 8. WG Grooming wing with rear legs; 9. MV Moving around, not grooming; 0. NO No grooming. Motionless.

Hierarchy of decisions

Dawkins & Dawkins (1973, 1974) analysed the temporal stream of events which is behaviour into a sequence of 'decisions'. A decision is defined as an event which itself could not easily be predicted, but from which future events can be predicted. Let letters of the alphabet represent behavioural events, and let the following be part of an observed sequence (ignore the underlinings for the moment):

VBQACVBQMFWACMFWACVBQACVBQMFWACAC
CACVBQACMFWVBQVBQMFWACVBQACMFWMFWAC

Analysis of transition frequencies shows the following: V is always followed by B which is always followed by Q. However, Q may be followed by a variety of events. Therefore, the triplet VBQ may be regarded as a unit. The 'decision' to do BQ is said to be taken at the same moment as the decision to do V. V is called a 'decision-point'. All decision-points in the above sequence are underlined. Since BQ is redundant following V, C is redundant following A, and FW is redundant following M, the above sequence can be more economically represented without loss of information, as a sequence of decisions: VAVMAMMAVAVMAAAVAMVVMAMMA. This advantage of parsimony is enjoyed not only by the ethologist seeking to reduce the volume of his data to manageable proportions, but also by the animal itself struggling to control its many muscles in efficient temporal patterns.

This is a hypothetical extreme example. Real-life decisions are not absolute but relative. We tried to measure, in bits of information, the relative 'decisioniness' of successive frames of film of chicks drinking. We found not just two sorts of event, unpredictable decisions and predictable follow-ups of decisions, but rather a smear of intermediates with perhaps local modes. More interestingly from the present point of view, we speculated that there might be a hierarchy of decisions in the following sense. There might be some predictability between successive *decisions*, for instance VBQ and AC in the above example might be more likely than chance to alternate with each other. Then in the sequence MFWVBQACVBQACVBQACMFW, the first V would constitute a bigger decision than subsequent Vs or As, because it signals the onset of a new VBQAC cluster. Subsequent Vs and As are more predictable, hence smaller decisions; they are within-cluster decisions. They are still called decisions because they are less predictable than the Bs, Qs and Cs. Similarly there might be an even bigger, more global decision to enter the whole MFWVBQAC major cluster, as opposed to some other major cluster involving acts not yet mentioned, say XYZ.

A model along these lines is specified in the following assumptions. Like

any non-trivial model this one is a piece of at least partially free invention which does not follow logically from known facts.

Assumptions of the model

The organisation of an animal's decisions is hierarchical if it is possible to group its behaviour patterns into clusters such that

(i) For each cluster there exists a state of the animal of being certain to do one element of the cluster, but still uncertain which. Thus A, B and C form a cluster if the animal is capable of entering a state in which it is definitely about to do (A or B or C) and nothing else, but this state still leaves open *which* of the three will be done (they are not necessarily equi-probable).

(ii) The elements within a cluster between which a choice is made may be single acts, or they may be subclusters defined in the same kind of way as the cluster under discussion. For example the animal might be capable of entering a state of being about to do (A or B or C or D or E) but nothing else. At other times it enters states of being certain to do (D or E) or of being certain to do (A or B or C), or of being certain to do D. However, there is no state of being certain to do (A or D) but not anything else, since A and D belong in different clusters at the same level; a state of being about to do A or D implies the possibility of doing B or C or E, the other members of the smallest cluster to which A and D both belong. The cluster of five acts can thus be represented as two subclusters:

((A or B or C) or (D or E)).

(iii) Choices may be influenced by previous choices only within clusters not between clusters, and only by previous choices during the current entry of the current cluster. Thus if a transition is observed between A and D this implies that the animal must have left Cluster 1 (A or B or C) and entered Cluster 2 (D or E). By assumption (iii), the choice of which member of the new cluster is performed is uninfluenced by which members of the old cluster had been chosen. Thus the transition A → D has a probability which is equal to the probability of the transition Cluster 1 → Cluster 2, with appropriate weighting for the overall rarity or commonness of A and D within their respective clusters. The same applies if A and D stand for subclusters rather than observed behaviour patterns.

A further assumption which is not essential but which it might be interesting to follow up is

(iv) Every decision the animal makes is a binary decision, a choice between two possibilities. This means that (A or B or C) must be dissectable into, for example ((A or B) or C). Any observed behaviour is then the consequence




Fig. 4. From top to bottom, sequence of decisions leading to final choice of behaviour B, as explained in text. Illustrates equivalence between tree diagrams and bracket notation. Each inverted T-junction is a choice point.

of a particular series of binary decisions. In the case of B the decisions are

- (((A or B) or C) or (D or E)) rather than anything else;
- ((A or B) or C) rather than (D or E);
- (A or B) rather than C;
- B rather than A.

Whether we adopt assumption (iv) or not, each successive decision is

represented in this notation as entering a more deeply nested pair of brackets. It could equally well be thought of as dropping one level in a tree-diagram until a terminal branch (actual behaviour) is reached (Fig. 4). The equivalence of these two notations is familiar from 'list processing' computer techniques (Foster, 1967), techniques which form the basis of most of the programmes to be described below.

Evidence bearing on the assumptions

Assumptions (i) and (ii) were in a sense part of an earlier model which was rigorously tested through its quantitative predictions (Dawkins, 1969a; Dawkins & Dawkins, 1974 and references cited therein). This was a model of choice between external stimuli, for example coloured spots between which a chick might choose to peck. If a chick prefers red to blue, and blue to green, the following 'states' were admitted by the model: choice of red alone; choice of red or blue; choice of red or blue or green. The model was originally expressed in terms of 'thresholds', but in the bracket notation it can be represented as (((red) or blue) or green). The extra brackets round red indicate that there is a state in which only red can be chosen. An additional strong assumption was made about choices being exactly equally distributed among all colours eligible according to the hierarchical decision rule at any given time. This assumption flowed naturally from the 'threshold' way of picturing the model, but it will not be discussed here since it is not relevant to the idea of hierarchy itself. It led to the possibility of strong quantitative predictions being precisely deduced and tested. The success of these predictions gives some confidence in assumptions (i) and (ii) of the present model.

Assumption (iii) implies that there is not just one global set of transition rules governing all behaviour patterns of an animal, which can most economically be expressed in a single transition matrix. Rather it postulates nested sets of transition rules, each set of rules holding sway within a circumscribed cluster of elements. Transitions between elements which belong in different clusters at any particular level are predictable from the more global rules of transition between their respective clusters.

Fentress and Stilwell (Fentress, 1972; Fentress & Stilwell, 1973) counted frequencies of transition between different components of face-grooming in mice, and used an information measure of predictability in transition matrices of various order. The sequential structure revealed by these analyses fell far short of what the unaided human eye seemed to see. They therefore defined five higher-order units, recognised by eye, and given semi-precise definitions such as: 'Unit 4, repeated overhands with rare short licking interjected'. When frequencies of transition between these higher-order units were examined, a

folgt auf		I							II									III								
		Si	Fr	Ny	La	Fl	11Bo	11Lu	11	112	1f1	11R	11T	11K	1A1A	1A1r	1A1	33	332	33H	33F	3H3F	33T	3T3H	3T3F	
I	Si		1	16	7	1	6	60	6		1					1		7		2	5					
	Fr	1		18	16	2	18	31	25			2			1	1		4	4							
	Ny	10	13		311	1	19	194	24	4	4	5			1	1	1	12	8	1	3			2	1	
	La	5	32	95		4	25	352	36	1	6	2			8	5	1	22	4	6	2				3	
	Fl					10		13																		
II	11Bo	8	19	50	14	2		110	41	3	2	3		1		5	3	11	1						1	
	11Lu	67	27	289	179	6	148		329	7	19	55	8	2	14	5	11	81	30	29	18		1	1	12	
	11	4	10	74	17		37	293		84	335	123	29	16	202	68	60	15	8	6	4					
	112		2	10	1		7	77		7		3		8				1	5							
	1f1		1	8	5		2	25	304	13		8	13	5	19	1	1		1	2						
III	11R		2	4	1		4	11	156		3			1	7	5	1									
	11T						4		41	3	5															
	11K			1			2		14		2				1	6	2									
	1A1A					2			101	1		5	1	1		113	57									
	1A1r			4	3	1			187	1	1	1		1	17		3									
III	1A1			1		2	4		124	1	1	1			1	3										
	33	14	4	36	10		3	147	13	3	1				1	2	3	80	111	162	4	18	13	138		
	332	1	1	8	2		21				1							109		3				1		
	33H	2		5	1		11											168	2	53	13	1		5		
	33F			2	3		13		2								1	268	2	68		5		6		
	3H3F																	1	2	1		1				
	33T		1				1											12	1	9	1		1	4		
	3T3H																	7	1	7				5		
	3T3F			1	2		2											44		29	83	2	2	4		

folgt		I							II										III						
auf		Si	Fr	Ny	La	Fl	11Bo	11Lu	11	112	1f1	11R	11T	11K	1A1A	1A1r	1A1	33	332	33H	33F	3H3F	33T	3T3H	3T3F
I	Si																								
	Fr	$\chi^2 = 844.6$							$\chi^2 = 59.1$										$\chi^2 = 34.4$						
	Ny	d.f. = 25							d.f. = 48										d.f. = 42						
	La	...							N.S.										N.S.						
	Fl	$C = 0.528$							$C = 0.290$										$C = 0.334$						
	11Bo																								
	11Lu																								
II	11																								
	112																								
	1f1	$\chi^2 = 37.5$							$\chi^2 = 1693.7$										$\chi^2 = 9.0$						
	11R	d.f. = 48							d.f. = 49										d.f. = 56						
	11T	N.S.							...										N.S.						
	11K	$C = 0.254$							$C = 0.656$										$C = 0.421$						
	1A1A																								
	1A1r																								
	1A1																								
III	33																								
	332																								
	33H	$\chi^2 = 10.9$							$\chi^2 = 0.9$										$\chi^2 = 730.5$						
	33F	d.f. = 42							d.f. = 56										d.f. = 36						
	3H3F	N.S.							N.S.										...						
	33T	$C = 0.190$							$C = 0.174$										$C = 0.579$						
	3T3H																								
	3T3F																								

Fig. 5. Contingency tests on portions of Seibt's (1972) table of transition frequencies, as explained in text. Upper table gives frequencies with which acts listed as row headings were followed by acts listed as column

high degree of predictability was found. They conclude, 'These data provide a direct demonstration of the hierarchical structure in ongoing grooming behaviour which has often been postulated by workers in both the behavioural and neurological sciences.'

In our fly grooming study already cited, we made a half-successful attempt to test assumption (iii) directly (Dawkins & Dawkins, 1976). Here I shall illustrate the same method using data from other published literature.

Myrberg's method of arranging his tables of transition frequencies into sub-tables has already been mentioned. Seibt (1972) gives a similar table of frequencies of transition between grooming and other movements in diopsid flies, which she divides into three groups. Group 1 are non-grooming movements, group 2 grooming movements involving the front legs, and group 3 grooming movements involving the rear legs. As in our study of blowflies she found that group 2 and group 3 grooming movements tended to be separated from each other, and she concluded that they formed two self-contained complexes, each with its own 'grooming-programme'. She separated group 1 movements on the same kind of basis.

As explained above, mere temporal clustering is not in itself evidence for hierarchical organisation in Simon's (1962) sense or in the sense of the present model. A direct test of assumption (iii) would consist of a demonstration that, in the case of between-cluster transitions, which member of the second cluster is chosen is not influenced by which member of the first cluster had previously been performed. In the case of within-cluster transitions, however, considerable sequential influence may occur. For Seibt's data for instance, in those cases where a transition is observed from a member of group 2 to a member of group 3, the probability that the latter will be any one of the eight members of group 3 should be uninfluenced by which of the nine members of group 2 the former was. This may be tested by subjecting the sub-table involving transitions from any member of group 2 to any member of group 3 to a 9×8 χ^2 contingency test. χ^2 for 56 d.f. is only 9.043 which

headings. Lower table shows results of contingency test on each of the nine main sections. Stars indicate statistical significance. In general prediction is fulfilled that significant interactions occur in submatrices around main diagonal, but not in other portions. To correct for effects of large numbers, Pearson's coefficient of contingency $C = \sqrt{\chi^2/(N + \chi^2)}$ is calculated for all nine cells, where N = total number of contingencies. Again prediction supported. Original table has no entries along main diagonal - behaviour patterns by definition cannot follow themselves. Strictly speaking this necessitates special treatment in calculating χ^2 (Slater, 1973). However, the effect is so large that the precaution can be omitted. N.S. = not significant; d.f. = degrees of freedom.

Preceding events	Following events															Row totals
	Flutter	Tilt	Dip	Nudge	Lead	Quiver	Close swim	Skim	To nest	Van	Nip eggs	Away from nest	Lawn	Tail down	Tail dig	
FL																237
TI																272
DI																2947
NU																378
LE																671
QU																343
CLS																62
SK																748
TN																2015
FA																1868
NIE																2334
AN																1989
YA																145
TDO																38
TDI																41
CHF																91
FF																29
NIR																182
NISW																7169
NIBW																1773
NISN																43
NISA																316
SO																239
FT																251
FD																206
LD																165
CHS																319
FS																111
HS																28
HU																106
HD																10
DU																162
DD																42
PS																42
TB																58
CI																45
HCI																12
	238	273	250	250	250	250	250	250	250	250	250	250	250	250	250	28387
	238	273	250	250	250	250	250	250	250	250	250	250	250	250	250	Total events

Fig. 6. Contingency tests on portions of Myrberg's (1972) Fig. 32. Row and column headings copied exactly from Myrberg. Three main submatrices around main diagonal are emphasised by a double line. All portions of the table not involving an interaction among these three ruled out (diagonal line). In remaining nine portions contingency tests were done as explained in text. Stars indicate statistical significance. In general prediction is fulfilled that significant interactions occur in the three square submatrices, but not in other portions (except one). To correct for effects of large numbers in square submatrices, Pearson's coefficient of contingency $C = \sqrt{\chi^2 / (N + \chi^2)}$ is calculated for all cells. Again prediction supported. To eliminate effects due to behaviour patterns leading to themselves, all entries along main diagonal of original table have been eliminated.

accords with the prediction. On the other hand if the members of group 2 are subjected to a within-cluster 9×9 contingency test, χ^2 is 1693.690, showing highly significant within-cluster interaction again as predicted. Fig. 5 gives the equivalent χ^2 for all nine major subdivisions of Seibt's table. The prediction that the three entries along the major diagonal may be highly significant but all the rest will not be, is fulfilled.

The row and column headings of Fig. 6 are reproduced directly from Myrberg's (1972) Fig. 32. The body of the table is divided into the main regions discernible in Myrberg's arrangement ('partial decomposability analysis' described above). The detailed figures are replaced by a χ^2 value and contingency coefficient for each region of the table. The prediction of the model is fulfilled, with the exception of one cell.

A cluster analysis based on Mutual Replaceability*

Normally a cluster analysis begins with a matrix of similarities or distances between elements (Everitt, 1974), and proceeds to group elements together which have high similarity. The single-linkage cluster analysis illustrated above (Fig. 3b) used temporal proximity as the index of similarity. Grooming acts were likely to be clustered together if they tended to occur in the same 5 second period. The method of rearranging transition matrices so as to maximise the concentration of entries in square submatrices is, as we saw, a form of two-levelled cluster analysis; in this case the index of similarity was probably of sequential contiguity.

But for the present model the appropriate index of 'clusteredness' of two behaviour patterns is not any form of temporal proximity, but rather *mutual replaceability* (cf. Kalmus, 1969) as far as between-cluster transitions are concerned. For example the blowfly head-grooming and proboscis-grooming movements belong in the same cluster, not because they tend to occur close to each other in time, but because their transition-relationships with members of other clusters are nearly the same: they are mutually substitutable in those parts of a transition matrix which do not involve their relationships with members of their own cluster. The χ^2 analysis above tested this kind of prediction using preconceptions about which behaviour patterns ought to be clustered together. The following method of cluster analysis was developed to discover which behaviour patterns are clustered together given no specific preconceptions other than the assumptions of the model itself. Assumption (iv), the one about all decisions being binary, was included for the mundane reason that it made computation easier.

* When I wrote this I was unaware that similar methods had been used before (Maurus & Pruscha, 1973).

The programme is provided with an ordinary first-order transition matrix as data. The arrangement of the rows and columns is irrelevant. It examines all possible pairs of behavioural acts in turn, calculating for each pair an 'index of mutual replaceability', which is the mean of two correlation (Spearman rank unless otherwise stated) coefficients r_r and r_c . r_r is the correlation between two rows, excluding the entries involving mutual interaction within the pair under investigation. r_c is the corresponding figure for the two columns. Having found that pair with the highest index of mutual replaceability, it designates them as members of the same cluster and prints out their names bound together in brackets. Following assumption (iii) of the model it then collapses the table so that no further distinction is made between these two behaviour patterns; their entries are lumped by simple addition. The whole operation is then repeated on the condensed table, and this continues until only two entries are left, or until no good correlation can be found, as defined by an arbitrary criterion. At each stage the pair of elements with the highest index of mutual replaceability is printed out; sometimes these two elements are single acts; sometimes they are already identified and lumped clusters, in which case a nested bracketing notation is used.

The method can be used on any transition-frequency data. Examples using published data are given in Fig. 7a and b. The clusterings of Myrberg's behaviour patterns can be compared with those which he arrived at (Fig. 6). Fig. 7c shows its application to blowfly grooming where it can be compared with a Single Linkage method already discussed. Do not expect that conventional methods of analysis based on temporal or sequential proximity will necessarily give similar clusterings to the Mutual Replaceability method. Two acts which are mutually replaceable would very probably *not* be sequentially close. To use Kalmus's (1969) analogy of a menu, two different fish dishes are mutually substitutable in the second slot of a four-course dinner; they are therefore unlikely to be served up in succession.

Fentress & Stilwell (1973) conclude by suggesting an analogy between mouse grooming and 'human grammar in which individual letters form different combinations in different words which in turn are sequentially arranged into phrases...'. As they point out similar suggestions have been made before. These will be discussed later. To anticipate, the main reason why I have preferred not to call the model discussed in this section grammatical is that a grammar is more than just a hierarchical system in which the higher units as well as the lower units have their own laws of transition. In addition, a sentence has a definite structure; at its crudest, it has a beginning, a middle and an end. A better way of putting this is in terms of 'correct nesting of brackets', and we will return to this later.

The behavioural analogy is with syntax not semantics, and it has nothing necessarily to do with the communicatory role of language. In some ways it

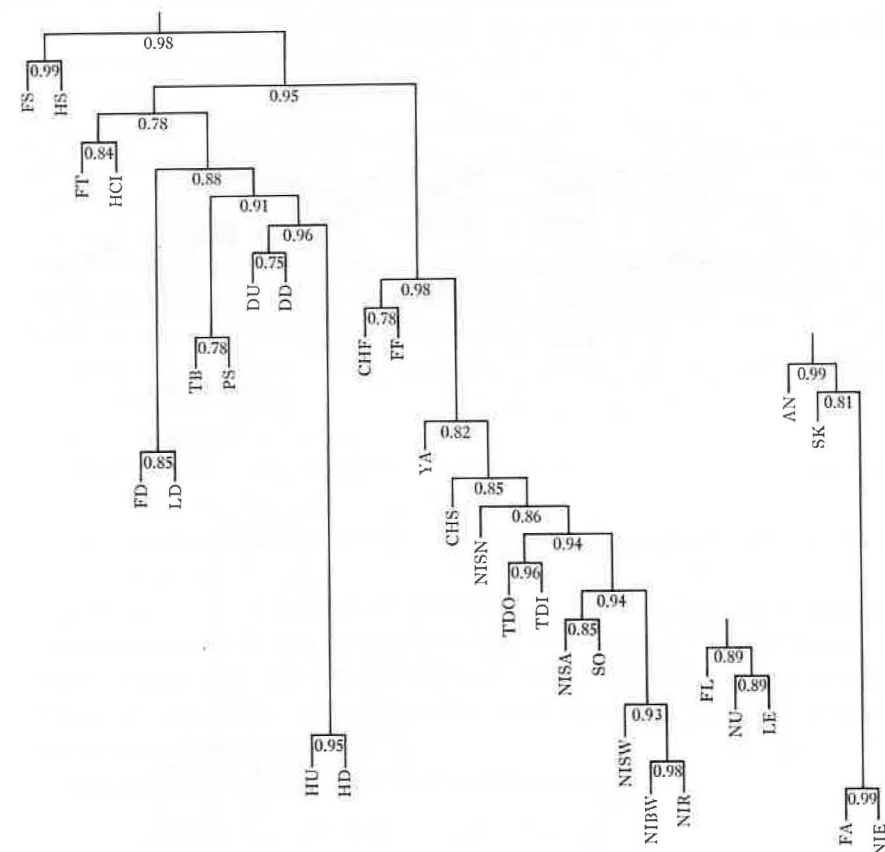


Fig. 7. (a) Mutual Replaceability Cluster Analysis on data from Myrberg (1972) on bicolour damselfish. Behaviour names as in Fig. 6. Each clustering of two subunits represented by one inverted T-junction. Members of earlier-formed clusters (*lower*), are lumped for later clustering (*higher*). Under each T-junction is Index of Replaceability for the two units joined. This is the mean of two correlation coefficients as explained in text. In this case product moment correlations used, as Spearman rank too costly of computer time.

is better to begin by comparing behaviour with simpler artificial 'languages' which were never designed for communication. These are the subject of the next section.

PATTERNS OF PATTERN

Several authors have developed what may be called pattern languages in the course of studies of serial pattern learning, the learning by human subjects of long sequences of symbols, digits, letters or responses. Simon (1972) has




Fig. 7. (b) Mutual Replaceability Cluster Analysis on data from Baerends, Brouwer & Waterbolk (1955), on courtship behaviour of male guppies, *Lebistes reticulatus*. As Fig. 7a, except Spearman rank correlations used. Upper part of figure: representation of transition frequencies by original authors. Behaviour symbols are those of original authors.

shown that, although the formulations of the various authors seem different, they are mostly equivalent. I take the work of Restle (1970) as representative. He used a box with six numbered buttons and six corresponding lights. The lights flashed on in a non-random repeated pattern. The subjects' task was to anticipate each light by pressing its button first. Some serial patterns proved harder to learn than others, and mistakes occurred at certain points in the sequence more than others. The aim was to explain these findings in terms of theories about how the subjects encoded the sequences in their memories.

Naive 'stimulus-response chain' theories could quickly be dismissed, including more sophisticated versions in which each response was associated not

only with its immediate predecessor, but with several predecessors. The subjects then were not behaving like a Markov process of any order. Nor were they simply storing the information in an ordered set of 'locations' as a computer might; if they were they would have learned a random sequence as readily as a patterned one. Restle used simple pattern languages to develop his theories of how they were doing it. These languages were ways of expressing a sequence using fewer symbols than the sequence itself. In each case the hypothesis was that the subjects were remembering, not the full sequence, but a set of *rules* for generating it, rules expressed in some formal equivalent of the 'language' under investigation.

A set of rules – I shall call it a programme – for generating a sequence would take the general form

- (i) start with button x ;
- (ii) perform some transform on x to select the next button, for instance move to the next one on the right;
- (iii) and subsequent steps. Perform some transform on the previous thing you did.

The important point is that in stage (iii), 'the previous thing' does not necessarily mean 'the last button'. The successful languages were all recursive, which means that 'previous thing' might be the last actual response, but it might be instead the last transform executed, either on a single response, or on an inner nested transform.

Particular transforms suggested by Restle were:

$T(x)$ move one to the right of x

$R(x)$ repeat x

$M(x)$ do the mirror of image of x

Thus $T(1) = 2$, $R(1) = 11$, $M(12) = 65$ since 6 is the mirror image of 1 on the button box, and 5 the mirror image of 2. $T(12) = 23$, $R(12) = 1212$, $R(R(12)) = 12121212$, $T(R(1)) = 1122$, $M(T(R(1))) = 11226655$. In all cases an equivalent tree-diagram can be drawn. For example the last 'programme' corresponds to Fig. 8a.

The sequence 11662255116622552255334422553344 is an example of one which proved easier to learn than one would naively suppose from its length. Restle suggested that this was because such sequences have a hierarchical structure, which in this case can be represented by $T(R(T(M(R(1))))))$. On the assumption that subjects represented the sequence internally by a set of rules equivalent to this formula, predictions were made about where in the sequence errors should be most likely to occur. They, and similar predictions for other sequences, were fulfilled. That particular 'programme' may be decoded into behaviour as follows: do behaviour 1; repeat it; do the mirror




Fig. 8. (a) Tree-diagram illustrating 'programme' $M(T(R(1)))$ for Restle's button box. (b) Equivalent diagram for hypothetical fly grooming programme $F(R(R(P(1))))$. Explanation in text.

image of the whole repeated unit; transpose both mirror halves; repeat everything done so far; transpose the whole repeated, transposed, mirrored, repeated unit.

Symmetrical binary trees have obvious limitations, for example they can only generate sequences whose number of elements is a power of two. Restle therefore also considers 'right-branching' asymmetrical trees. Using them he makes a spirited attempt to analyse J. S. Bach's Two-part Invention No. 1, and he suggests that the method is applicable to many forms of human behaviour, including such complex skills as playing the piano. I now leave human psychology, and return to animal behaviour.

The nature of the proposed analogy should be clear. The animal is supposed to generate sequences of behaviour by following economically stored sets of rules of the same type as those used by the subjects in the learning experiments. The following is a sequence of behaviour recorded from a grooming blowfly: 1313131313131414167676 (Dawkins & Dawkins, 1976). Each digit stands for one act, for example 1 is rubbing the front legs together, 3 is grooming the head with the front legs, 6 is rubbing the back legs together and 7 is grooming the abdomen. The most obvious patterns that can be seen are what Restle would call 'trills', long periods of alternation between say 1 and 3, and 1 and 4. However, before we rush into a direct search for

hierarchical rules we must be careful. The numbers on Restle's box had some ordinal significance; thus a 121212 trill represented repeated pressing of a button and its neighbour. But the assignment of numbers to grooming movements is purely arbitrary. The operations 'transpose' and 'mirror' have no obvious meaning.

This is not to say that some sort of meaning for them could not be found. For example we might hypothesise the following transformations, some of them making use of the concept of 'postural facilitation' (Dawkins & Dawkins, 1976):

- $R(x)$ As in the case of Restle, repeat x
- $A(x)$ Groom that part of the body immediately anterior to x
- $P(x)$ Groom that part of the body immediately posterior to x
- $M(x)$ Groom that part of the body which is the left/right mirror image of x
- $F(x)$ Groom that part of the body which is the fore and aft mirror image of x , for instance 'rub rear legs together' might be the mirror image of 'rub front legs together'

As before, x might refer to a single grooming act, or more interestingly it might refer to a higher order unit as in $F(R(R(P(\text{front leg rubbing}))))$. This could be decoded as: rub the front legs together; move one stage posteriorly and groom the head; repeat the whole thing twice; then do everything again but fore/aft reversed, substituting back legs for front legs and abdomen for head (Fig. 8b). The whole sequence would be 1313131367676767. Predictions could be made and tested about which sorts of sequences should and should not occur commonly if fly behaviour is generated by various programmes of calls of these procedures. We have not tried this yet, but it would be interesting to do so, and there may be other kinds of animal behaviour, bird song perhaps (cf. Nelson, 1973), for which a similar approach could be worthwhile. This is a suggestion for the future.

An algorithm for detecting patterns in behaviour - 'melodies'

The suggested transpositions given above are based on human preconceptions of what might be reasonable. Another approach is to scan the data for patterns which actually do occur, and for patterns of pattern. The following algorithm attempts to do this.

The programme is provided with raw data consisting of digits signifying acts, in the order in which they occurred. It scans through the data counting frequencies of doublet transitions. When it has found the commonest doublet it prints out the two behaviour names in order, bound together in brackets.

Frequency	Doublet
+83	(6 0)
+79	(1 3)
+65	((1 3)(1 3))
+51	(7 (6 0))
+45	(1 0)
+25	(1 2)
+23	((((1 3)(1 3))((1 3)(1 3))))
+21	(7 6)
+19	(1 4)
+16	((7 (6 0))8)
+15	((6 0)8)
+14	((1 2)(1 2))
+13	((1 0)(1 0))
+12	((1 4)(1 4))
+11	((7 (6 8))6)
+9	((7 (6 0))(7 (6 0)))
+7	(6 8)
+6	(5 (6 0))
+6	((1 3)(1 0))
+5	(0 (6 0))
+5	(6 (7 6))
+5	((1 0)(1 2))
+5	((((1 3)(1 3))((1 3)(1 3)))((1 3)(1 3)))
+5	((7 6)(7 6))
+5	((((1 2)(1 2))((1 2)(1 2))))
+4	(6 ((7 (6 0))6))
+4	(7 0)
+4	((1 0)4)
+4	((((6 0)8)((6 0)8)))
+4	((((1 0)(1 0))(1 0)))
+3	((7 (6 0))(7 6))
+3	(((((1 3)(1 3))((1 3)(1 3)))((1 3)(1 0))))
+3	((((7 (6 0))8)(7 (6 0))))
+3	((((7 (6 0))8)((6 0)8)))
+3	((((1 4)(1 4))(1 4)).

Fig. 9. Computer recognition of patterns in sequences of grooming movements by one blowfly 'April' (from Dawkins & Dawkins, 1976). Common doublet sequences of acts, and doublet sequences of already detected doublets (etc.), printed out in order of commonness. e.g. doublet 6→0 occurred 83 times and was commonest. When this was replaced by a single symbol (6 0), commonest doublet in altered record was 1→3. When this was replaced by (1 3), commonest remaining doublet was (1 3)→(1 3), which occurred 65 times, and so on until no doublets which occurred more than twice could be found. The two halves of each doublet may be identified from nesting of brackets.




Fig. 10. Illustrations of some of the commoner 'doublets' from Fig. 9. Data from grooming sequences of a blowfly. Numbers refer to grooming acts as listed in caption to Fig. 3.

It then goes back over the data, replacing all occurrences of that doublet by a single symbol which stands for the doublet. It then repeats the process on the altered record. At each stage, one doublet is printed out, the elements of the doublet being either single behaviour names, or already identified patterns, represented in a nested bracket notation. Similar algorithms have been developed for detecting the natural segmentation of language (Wolff, 1975).

Superficially this programme may seem like the other one described above,

but they are doing quite different things. The first one used a table of doublet transition frequencies as its data, and it clustered pairs of acts (or already formed clusters) on the basis of their mutual replaceability. This programme on the other hand uses raw sequence data, and progressively builds up a picture of the commonest sequential patterns, including patterns made up of already detected patterns.

These common patterns might be called melodies. Indeed, as in the case of mouse grooming (J. C. Fentress, personal communication), if each blowfly grooming act is represented by a musical note played by a computer, the human ear seems to pick out distinctive melodies. The effect is rather like 'modern' jazz, and there are resemblances between different individual flies. Fig. 9 shows the common 'melodies' in the order (decreasing frequency) in which they were extracted from data by the programme. The results are also expressed in Fig. 10 in diagrammatic form.

Another type of hierarchical model, which also has affinities with grammatical models, is based on the idea of a 'goal'.

HIERARCHY OF GOALS

Until recently 'goal-directed behaviour' was claimed as a profound mystery by those who enjoy profound mysteries (e.g. Russell, 1946). Biology does still present mysteries, but goal-directed behaviour is not one of them. It was reduced to the commonplace, not by philosophy but by gunnery (Rosenblueth, Wiener & Bigelow, 1943). There is still the possibility of confusion between goal in the cybernetic sense of 'stopping condition' (Hinde & Stevenson, 1970), and goal in the sense of survival value. The latter sense is misleading and should not be used.

The idea of a hierarchy of goals is well known from Miller *et al.*'s (1960) book *Plans and the Structure of Behavior* (see also Pribram, 1971, in which the idea is extended by the incorporation of feedforward). They are mainly concerned with human psychology, but their book has a chapter which calls attention to the similarity between their own hierarchical 'Plans', and the hierarchical 'Instincts' of Tinbergen, which they discuss with approval. They might have been even more pleased by Kortlandt's work (1955) in which, although it can be criticised (Hinde, 1957), the link with their own ideas was more direct.

Kortlandt's is a hierarchy of 'appetites', where an appetite is defined by the condition which brings it to an end. He saw the nest-building behaviour of his cormorants as mediated through the arousal of hierarchically subordinate appetites. Fig. 11 illustrates his term 'concentric purposiveness' in two equivalent ways. It could also be easily drawn by Miller *et al.* as a 'TOTE'

HIERARCHICAL ORGANISATION

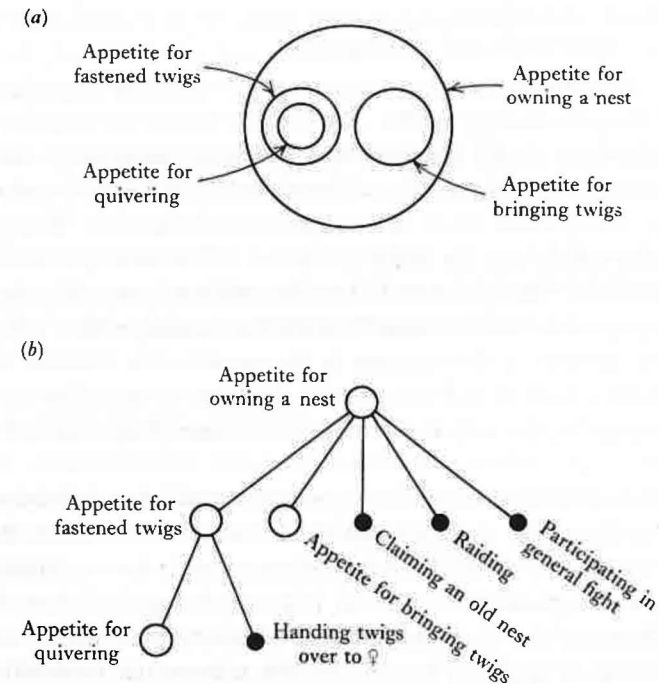


Fig. 11. 'Hierarchy of appetites' in the male cormorant *Phalacrocorax carbo* (from Kortlandt, 1955). Each appetite is defined by the state which brings it to an end. It may achieve its 'goal' by calling up subsidiary appetites. Thus the appetite for owning a nest is superior to the appetite for fastened twigs which in turn is superior to the appetite for quivering. (a) 'Concentric purposiveness'. (b) Equivalent, but extended, tree-diagram.

hierarchy. Thorpe (1963) gives an illuminating discussion of how nest-building and artifact construction in general can be understood in similar terms.

Hinde and Stevenson (1970), also in the context of nest-building, warn against over-enthusiastic interpretation of behaviour in terms of goals. They go on to provide an admirable classification of the rather diverse kinds of phenomena which might be called goal-directed, and suggest that for classification purposes the category may be too heterogeneous to be useful. Nevertheless I believe the following general functional or economic argument may be worth putting, since it lays particular stress on the hierarchical organisation of goals. It has affinities with Simon's watchmakers.

Action rules and stopping rules

We may distinguish two extreme strategies for programming adaptive behaviour. The goal strategy is this. The animal knows nothing about how to behave adaptively except a *stopping rule*. The programme simply says 'Thrash around at random until state G is achieved, then stop'. Orthokinesis (Fraenkel & Gunn, 1940) is not much more sophisticated than this. That part of the programme which may be called the *action rule* is as simple as it could be, 'thrash around'. The goal state G may be simple as in orthokinesis, or it may be very complex. Under the name 'British Museum Algorithm' a slightly more systematic version of this strategy is discussed in the artificial intelligence literature as a kind of null hypothesis. In theory it can solve any problem, however complex, but only at enormous cost in time (Feigenbaum & Feldman, 1963).

The opposite extreme would be a programme which put all the complexity into the action rules, a programme of the form 'Do A then do B then if X do C otherwise...' This is the type of programme conventionally (though not necessarily) given to computers. It is very fast and efficient, but only if the environment can be relied upon to be utterly predictable; otherwise it will fail when it stubs its toe on the first unexpected mole-hill. It is not possible to plan in detail for every contingency; there are too many of them.

Is there any way of combining the virtue of stopping rule programmes, imperturbability in the face of unpredictable conditions, with the virtue of action rule programmes, speed and efficiency? Yes. The solution lies in hierarchically nested stopping-rule programmes. This is the basis of the models of Kortlandt, and of Miller *et al.*

Stopping-rule programmes can be fast, provided the goal state is simple. This may be because it is sometimes possible to measure quantitatively the discrepancy between the goal and the present state, in which case the full power of negative feedback can be brought to bear (McFarland, 1971). But even if this is not so, simple stopping conditions may be achieved rapidly, because simple means not improbable. A complex goal state, like the particular permutation of letters which is any book in the British Museum, is inherently improbable. A goal state like 'stomach full of zebra meat' is too complex to achieve through random movements. A predator programmed with only such a stopping rule might take millions of years to achieve a square meal. It is obviously better to break down complex and improbable goals, into a series of simple goals which can be more rapidly achieved. For instance immediately subordinate subprogrammes might be: stop searching when zebra seen; stop pursuit when zebra very close; stop killing when zebra motionless; stop eating when stomach full. 'Searching', 'pursuit', 'killing' and 'eating' are deliber-

ately left vague. Each of them would have its own programme, which might consist purely of action rules, but more probably, since they are all still quite complex, each one would call up its own subordinate stopping-rule programmes. Even at the very lowest level, there seem to be stopping-rule programmes, in the form of the γ -efferent servo loops. Action rules are perhaps mainly confined to determining the order in which stopping rules or targets are set up.

Another aspect of hierarchies of goals may be expressed as 'perfection of nesting', using the word in the sense of nesting of brackets! This is best discussed after considering grammatical models of behaviour.

GRAMMATICAL MODELS

The idea of some similarity between the principles underlying language and those underlying the serial organisation of behaviour in general is obviously of great interest from many points of view, including that of the evolution of language. It seems to have originated from Lashley (1951), but to have been first turned into an explicit model by Marshall in an unpublished paper in 1965. His 'phrase structure grammar' model did not become widely known until it was discussed in print by Hutt & Hutt (1970), and by Vowles (1970) who added some speculations about neurophysiological implications and about 'transformational grammar' (Chomsky, 1957). Meanwhile Kalmus (1969) had independently developed the analogy.

A major aim of grammarians at the time of Marshall's paper was to write down rules of the syntax of a particular language, in such a way that in theory a machine embodying these rules could generate all grammatical sentences recognised as correct by native speakers, and no ungrammatical sentences (Chomsky, 1957). Such a machine would only be concerned with syntactics, with deciding when to emit a noun, an adjective, a relative pronoun etc. The selection of *which* noun and which adjective is a semantic matter. It is not obvious what an analogy of semantics with animal behaviour would mean; perhaps something to do with the functional achievements of behaviour.

A computer programmed to follow some typical hierarchical and recursive rules of syntax, produced the following sentence (ignore the underlinings for the moment): 'The adjective noun of the adjective noun which adverbly adverbly verbed in noun of the noun which verbed adverbly verbed.' Dissect it carefully, and you will find that, although meaningless, it is syntactically correct English. The fundamental sentence is underlined: 'The adjective noun adverbly verbed'. All the rest consists of qualifications of the subject *noun*, in the form of hierarchically nested relative, prepositional and possessive clauses. The programme employed random numbers to make its successive


choices, and it continued to generate prolific grammatical gibberish. However, the interest of it is not in its randomness, but in the nature of the *units* between which each random choice was made.

If the units had been *words*, with choices biased to follow transition probabilities actually found in English, the sentences generated would at first sight have looked acceptable. Adjectives would precede nouns and adverbs would precede verbs. If the Markov process simulated was one of high order, whole phrases would come out looking plausible. But there would be one thing fundamentally wrong: clauses which had been begun would not be properly finished. If we represent the opening of a new relative or possessive clause as equivalent to the opening of a bracket, the Markovian programme would not correctly close all brackets opened. In terms of the above example, the Markovian model would drift off into relative and possessive clauses, and would not 'remember' that there was an initial subject still waiting for its main verb.* Even the 'decision cluster' model discussed above would not perform any better.

Our phrase structure grammar programme on the other hand achieved this correct rounding off of 'brackets' effortlessly, and it would have done so no matter how many subsidiary clauses had been opened. This is because each random number determined the choice not of a word, nor of a fixed number of words, nor even of a cluster from which words might be chosen, but of a *procedure*. These procedures, with names like 'noun-phrase' and 'verb-phrase' themselves used random numbers to choose either words or other procedures (including themselves) and so on. For example the procedure noun-phrase might generate any of the following: 'noun', 'adjective noun', 'noun of' noun-phrase, 'noun which' verb-phrase, and many others. Since the procedures noun-phrase and verb-phrase are recursive (i.e. they call themselves), there is no limit to the depth of nesting of 'brackets' which will be correctly rounded off.

Marshall (I am referring to secondary sources only) proposed a grammar of this type to generate the sequences of pigeon courtship described by Fabricius & Jansson (1963). The 'words' are seven behaviour patterns with names like Bow and Copulate. The higher units (equivalent to noun-phrase, etc.) have names like Preparatory and Consummatory. The complete grammar is given in Fig. 12, both in tree-diagram form, and in Algol-60. The latter is preferred to Marshall's own metalinguistic symbolism, which it closely resembles, as it may be directly run on a computer as a simulation of a pigeon, and the results of one such run are given in the Figure. It should

HIERARCHICAL ORGANISATION



```

begin comment Marshall's pigeon grammar;
procedure SBSeq; begin Prep; Con end;
procedure Prep; begin Int; Wa; if p then Prep end;
procedure Int; begin "BW"; if p then Agg; if p then Int end;
procedure Agg; begin if p then "DR"; if p then "A"; if p then Agg
end;
procedure Wa; begin "D"; "BI"; if p then Wa end;
procedure Con; begin "M"; "CO" end;
Boolean procedure p;
begin comment true or false at random. Probability manipulated.
end;
start: SBSeq; goto start
end of pigeon grammar;

```

Sample results of running the programme:

```

BW DR D M CO
BW A D BI BW DR D BW A D BW A D BI M CO
BW A D BI M CO
BW DR D BW DR D BI BW DR D BI BW A D BW A D M CO

```

Fig. 12. Marshall's pigeon grammar (recursive version) (from Hutt & Hutt, 1970, courtesy of C. Thomas Publ.) after Marshall (unpublished). Underneath is the same grammar written out in Algol-60, followed by some results of running the programme as a simulation of a pigeon.

be self-explanatory even to those who do not know the language (**begin** and **end** are brackets: the programme consists mostly of definitions of **procedures** in terms of other procedures, followed by repeated final 'calls' of the most global procedure **SBSeq**).

Marshall shows that his grammar accounts for many of the findings of Fabricius & Jansson. However, we have to remember that he chose one

* Compare the case of a woman with a severe lesion of the left frontal lobe, who wrote: 'Dear Professor, I want to tell you that I want to tell you that I want to tell you that I want to tell you...' and so on for several pages (Luria, 1970).

grammar out of a very large number of possible ones.* Moreover there is no indication that his grammar accounts for the data better than a Markov model, though it certainly does so more elegantly. With data predigested into the form of a transition probability table it is doubtful if a convincing test of the two models could be accomplished. Raw data or data processed as in Fig. 9 ('tunes') would probably be better, but I have not so far succeeded in framing definite testable predictions characteristic of grammatical models, and I can only put it as a challenge for the future.

The problem is this. In the case of human language the criterion for grammatical correctness is the judgement of a native speaker of the language. In the case of animal behaviour 'correct' and 'incorrect' have no such meaning (Altmann, 1965). It is not obvious for example what might constitute correct 'rounding off of brackets'. There seem to be two main ways of starting to look at it and it is best not to muddle them together. The first is in terms of pattern as discussed in the section on pattern languages. The second is in terms of goals as discussed in the previous section.

A 'correctly rounded off' serial pattern might show itself in observed behaviour in the following general way. Let A and B be two particular acts. Then we might observe that the sequence AXB is very common, where X is not one particular act, but any one of a set. In a sense then the performance of A opens a sequence which demands to be 'closed' by the occurrence of B, regardless of what X happens to be. Now let X stand for, not one single act but a series of unfixed length – a subsidiary pattern in fact. Thus we might see A12B, A12121212B and so on. In all of these the sequence AXB is still there, but no Markov chain analysis would ever detect it. It is a task for the future to find convincing examples of this, but a start has been made by Dawkins & Dawkins (1976).

Correct 'rounding-off' also has meaning with respect to goals. When a subsidiary goal is set up in the service of a more global one, the global one presumably can be still 'there' in whatever sense a goal is ever 'there'. Thus if a hyena sets up a subsidiary goal of 'the other side of the hill' in the service of the more global goal of 'catch zebra', it may be that in some sense the goal 'zebra' is still 'set up'. In the case of spotted hyenas (*Crocuta crocuta*) this is quite plausible, since Kruuk (1972, and personal communication) could tell in advance what prey hyenas were setting out to hunt, and they would not be distracted by the 'wrong' prey, even of a species which on another occasion they might set out to hunt. Correct 'rounding off' of goals then might

* Harding (1971) shows that the number of labelled binary tree shapes of degree n is $\frac{(2n-1)!}{2^{n-1}(n-1)!}$.

For seven types of pigeon behaviour, $n = 7$ and the number of possible grammars is 10395 discounting variants due to optional choices and recursion. The real number is very much larger, but many can be immediately rejected.

show itself as a tendency to return to a global appetitive pattern after subsidiary patterns have been initiated and satisfied, and also after distraction or interruption. Searching behaviour and the notion of a 'searching image' is of great interest to ethologists and ecologists (Dawkins, 1971; Krebs, 1973). The possibility that whatever is searched for may be hierarchically nested should be borne in mind in future.

If it is ever possible to attach meaning to 'correct' and 'incorrect' in animal behaviour, either in terms of statistical rarity, or in terms of functional appropriateness ('displacement activities?'), it might then be possible to analyse such 'mistakes' as do occur in the same kind of way as linguists have analysed speech errors to dissect underlying organisation (Fromkin, 1973). Spoonerism in song-birds would make a good subject for this centre of bird-song research (see Thorpe & Hall-Craggs, this volume).

SUMMARY COMPARISON OF BEHAVIOURAL MODELS DISCUSSED

Markov models of any order (our non-hierarchical 'null-models') have the following property. Events influence future events, and the degree of influence is less for the distant future than for the near future, a decreasing monotonic function of time. If behaviour sequences are regarded as the consequence of a corresponding sequence of decisions (as defined above), in the case of simple Markov models there is a one to one (or at least one to some fixed number) relationship between decision and observed act. In the case of all the hierarchical models this is not so.

In the 'decision cluster' model, the animal has to decide which cluster to enter, i.e. from which group of acts the final selection will be made. Then, after a series of subdecisions, it arrives at something the observer can actually see. As so far expressed, this is not different from a simple Markov model; it only sounds different. The crucial difference results from assumption (iii), that decisions in any cluster are not influenced by previous decisions within different clusters. This means that the decay of influence of present events on future events is not a constant function of time. It depends on whether a new cluster is chosen, in which case the decay is abrupt. Influence is still always less on the distant future than on the nearer future.

In the grammatical and 'pattern language' and goal models, not only is the number of acts per decision variable, but decisions about the distant future may be taken before decisions about the nearer future. This was expressed by the metaphor about brackets once opened having to be closed.

IS IT BIG ENOUGH FOR THE JOB?

At school we learnt the evidence for the theory of evolution, fossils, geographical distribution, hierarchical taxonomy and so on. Of course the evidence is very important – why otherwise should we remember Darwin more than Wallace? – but I confess it was not evidence that convinced me! The compelling thing about the theory of evolution is that it is big enough to do the job of explaining the otherwise inexplicable fact of our existence.

The nervous system will provide the last of the deep problems of biology. Anybody who has thought about it must be awed by what his own brain can do. Even its lesser accomplishments, the control of complex behaviour, the analysis of complex sensory data, the storage and rapid retrieval of voluminous memories, raise difficult enough problems. Perhaps their solution will never be anything but a mess of detail, but big problems invite big solutions. The question of how complexity in the world could come out of simplicity was answered in two words: natural selection. If we were forced to look into the future and guess which two words of our present meagre vocabulary might come closest to playing the same role for the understanding of complex behaviour, what would they be? Negative feedback? Powerful, but only in explaining simple behaviour: not big enough for the job. Hierarchical organisation?

I have emphasised the distinction between hierarchies of classification and hierarchies of connection. However, we have seen that some muddle between them is pardonable, and we can now see that the two may eventually come together. I think something like the principle we know as hierarchical organisation may turn out to be 'big enough for the job'. If so it will be for the same reason that it is an indispensable classificatory device. Whether it is complexity of stored information, complexity of pattern in incoming data, or complexity of controlled output, hierarchical organisation provides a way of making complexity manageable.

SUMMARY

(1) In the long term, general principles of Software Explanation of behaviour will be required. Hierarchical Organisation became unfashionable in ethology for the wrong reason.

(2) A hierarchy is defined semi-rigorously as a set of elements together with a relation called 'is boss of'. Linear, branching and overlapping hierarchies are also defined in the same way, and a distinction made between hierarchies of classification and hierarchies of connection. Possible examples of non-hierarchies are discussed.

(3) Three functional or 'neuroeconomic' arguments for regarding hierarchical design as good design are put forward. These are the Evolutionary Stability, the Local Administration, and the Redundancy Reduction advantages.

(4) In spite of statements common in the ethology literature, clustering of behavioural acts in time is not evidence for underlying hierarchical organisation.

(5) A truly hierarchical model is developed, based on the idea of 'decision' taken from previous papers. Animals are supposed to take global decisions initially, and to take progressively narrower subdecisions, ending up with an observed act. Published data are re-analysed to provide evidence bearing on the model. A method of cluster analysis based on the model, and called Mutual Replaceability Cluster Analysis is described, and tested.

(6) Special purpose 'languages' for describing serial patterns economically have been developed by human psychologists. These may be applicable to animal behaviour if we suppose that animals generate patterns by executing economical stored programs using hierarchical and recursive procedure calls. A hypothetical example is given, involving postural facilitation in grooming. An algorithm for extracting 'melodies', simple patterns, and patterns within patterns, is described and tested.

(7) Hierarchical models based on the idea of 'goal' are discussed. A distinction is made between two strategies of programming behaviour, based on 'action rules' and 'stopping rules'. Both have their disadvantages, and a good compromise is a set of hierarchically nested stopping rules.

(8) Models comparing animal behaviour with human syntax are discussed. An analogy with 'correct nesting of brackets' is developed. Recursive hierarchical rules of syntax have the interesting property that they effortlessly 'round off' main clauses even after subsidiary clauses, nested to an indefinite depth, have been initiated and concluded.

(9) The behavioural models used in the paper are compared according to two criteria: whether the number of acts per 'decision' is fixed or variable; and whether behaviour in the near future is necessarily easier to predict than behaviour in the more distant future.

(10) It is suggested that hierarchical organisation may be a generally powerful explanatory concept.

Many of the ideas in this paper were developed in discussions, over a long period, with Marian Dawkins. I am very grateful to her and also, for various reasons, to Bruce Anderson, Michael Arbib, Ted Burk, David McFarland, Sir Peter Medawar, Pat Searle, Richard Sibly, Wilson Sutherland, Niko Tinbergen, and the Editors of this book.

REFERENCES

- Altmann, S. A. (1965). Sociobiology of rhesus monkeys. II. Stochastics of social communication. *Journal of theoretical Biology*, **8**, 490-622.
- Attneave, F. (1954). Informational aspects of visual perception. *Psychological Review*, **61**, 183-193.
- Baerends, G. P. (1941). Fortpflanzungsverhalten und Orientierung der *Ammophila campestris* Jur. *Tijdschrift voor entomologie* **84**, 68-275.
- Baerends, G. P., Brouwer, R. & Waterbolk, H. T. (1955). Ethological studies on *Lebistes reticulatus* (Peters). 1. An analysis of the male courtship pattern. *Behaviour*, **8**, 249-334.
- Barlow, H. B. (1961a). The coding of sensory messages. In *Current Problems in Animal Behaviour*, ed. W. H. Thorpe & O. L. Zangwill. Cambridge University Press: London.
- Barlow, H. B. (1961b). Possible principles underlying the transformations of sensory messages. In *Sensory Communication*, ed. W. A. Rosenblith. MIT Press: Cambridge, Mass.
- Barlow, H. B. (1972). Single units and sensation: a neuron doctrine for perceptual psychology? *Perception*, **1**, 371-394.
- Barlow, H. B., Narasimhan, R. & Rosenfeld, A. (1972). Visual pattern analysis in machines and animals. *Science, Washington*, **177**, 567-575.
- Bastin, E. W. (1969). A general property of hierarchies. In *Towards a Theoretical Biology*, vol. 2, *Sketches*, ed. C. H. Waddington. Edinburgh University Press: Edinburgh.
- Bullock, T. H. (1961). The problem of recognition in an analyser made of neurons. In *Sensory Communication*, ed. W. A. Rosenblith. MIT Press: Cambridge, Mass.
- Bullock, T. H. & Horridge, G. A. (1965). *Structure and function in the Nervous System of Invertebrates*. MIT Press: Cambridge, Mass.
- Chomsky, N. (1957). *Syntactic Structures*. Mouton: The Hague.
- Craik, K. J. W. (1943). *The Nature of Explanation*. Cambridge University Press: London.
- Davis, W. J., Mpitsos, G. J. & Pinneo, J. M. (1974). The behavioral hierarchy of the mollusk *Pleurobranchaea*. 1. The dominant position of the feeding behaviour. *Journal of Comparative Physiology*, **90**, 207-224. II. Hormonal suppression of feeding associated with egg-laying. *Journal of Comparative Physiology*, **90**, 225-243.
- Dawkins, M. (1971). Perceptual changes in chicks: another look at the 'search image' concept. *Animal Behaviour*, **19**, 566-574.
- Dawkins, R. (1969a). A threshold model of choice behaviour. *Animal Behaviour*, **17**, 120-133.
- Dawkins, R. (1969b). The attention threshold model. *Animal Behaviour*, **17**, 134-141.
- Dawkins, R. & Dawkins, M. (1973). Decisions and the uncertainty of behaviour. *Behaviour*, **45**, 83-103.

- Dawkins, M. & Dawkins, R. (1974). Some descriptive and explanatory stochastic models of decision-making. In *Motivational Control Systems Analysis*, ed. D. J. McFarland. Academic Press: London.
- Dawkins, R. & Dawkins, M. (1976). Hierarchical organisation and postural facilitation: rules for grooming in flies. *Animal Behaviour*, in press. **24**, 739-755.
- Deutsch, J. A. (1960). *The Structural Basis of Behavior*. Chicago University Press: Chicago.
- Dorsett, D. A., Willows, A. O. D. & Hoyle, G. (1973). The neuronal basis of behaviour in *Tritonia*. IV. The central origin of a fixed action demonstrated in the isolated brain. *Journal of Neurobiology*, **4**, 287-300.
- Eccles, J. C. (1975). In *The Creative Process in Science and Medicine*, ed. H. A. Krebs & J. H. Shelley. Excerpta Medica: Amsterdam.
- Eibl-Eibesfeldt, I. (1975). *Ethology. The Biology of Behavior*, 2nd edn. Holt, Rinehart & Winston: New York.
- Elsner, N. (1973). The central nervous control of courtship behaviour in the grasshopper *Gomphocerippus rufus* L. (Orthoptera: Acrididae). In *Neurobiology of Invertebrates*, ed. J. Salanki. Tihany: Budapest.
- Everitt, B. (1974). *Cluster Analysis*. Heinemann: London.
- Fabricius, E. & Jansson, A.-M. (1963). Laboratory observation on the reproductive behaviour of the pigeon. *Animal Behaviour*, **11**, 534-547.
- Feigenbaum, E. A. & Feldman, J. (1963). *Computers and Thought*. McGraw-Hill: New York.
- Fentress, J. C. (1972). Development and patterning of movement sequences in inbred mice. In *The Biology of Behavior*, ed. J. A. Kiger. Oregon University Press: Eugene.
- Fentress, J. C. & Stilwell, F. P. (1973). Grammar of a movement sequence in inbred mice. *Nature, London*, **244**, 52-53.
- Foster, J. M. (1967). *List Processing*. Macdonald/Elsevier: London.
- Fraenkel, G. S. & Gunn, D. L. (1940). *The Orientation of Animals*. Oxford University Press: Oxford.
- Fromkin, V. A. (1973). Slips of the tongue. *Scientific American*, **229** (6), 110-117.
- Gower, J. C. & Ross, G. J. S. (1969). Minimum spanning trees and single linkage cluster analysis. *Applied Statistics*, **18**, 54-64.
- Harding, E. F. (1971). The probabilities of rooted tree-shapes generated by random bifurcation. *Advances in applied Probability*, **3**, 44-77.
- Hinde, R. A. (1952). The behaviour of the great tit and some other related species. *Behaviour Supplement*, **2**, 1-201.
- Hinde, R. A. (1953). Appetitive behaviour, consummatory act, and the hierarchical organisation of behaviour - with special reference to the great tit (*Parus major*). *Behaviour*, **5**, 189-224.
- Hinde, R. A. (1956). Ethological models and the concept of drive. *British Journal for the Philosophy of Science*, **6**, 321-331.
- Hinde, R. A. (1957). Consequences and goals. Some issues raised by Dr. A. Kortlandt's paper on 'Aspects and prospects of the concept of instinct'. *British Journal of Animal Behaviour*, **5**, 116-118.

- Hinde, R. A. (1960). Energy models of motivation. *Symposia of the Society for experimental Biology*, **14**, 199–213.
- Hinde, R. A. & Stevenson, J. G. (1970). Goals and response control. In *Development and Evolution of Behavior*, ed. L. R. Aronson, E. Tobach, D. S. Lehrman & J. S. Rosenblatt. Freeman: San Francisco.
- Hodos, W. & Campbell, C. B. G. (1969). Scala naturae: why there is no theory in comparative psychology. *Psychological Review*, **76**, 337–350.
- Holst, E. von & Saint Paul, U. von (1963). On the functional organisation of drives. *Animal Behaviour*, **11**, 1–20.
- Hutt, S. J. & Hutt, C. (1970). *Direct Observation and Measurement of Behavior*. Thomas: Springfield, Ill.
- Iersel, J. J. A. van & Bol, A. C. (1957). Preening of two tern species. A study on displacement activities. *Behaviour*, **13**, 1–87.
- Ito, N. (1974). The control mechanism of cerebellar motor systems. In *The Neurosciences – Third Study Program*, ed. F. O. Schmitt & F. G. Worden. MIT: Cambridge, Mass.
- Kalmus, H. (1969). Animal behaviour and theories of games and of language. *Animal Behaviour*, **17**, 607–617.
- Kennedy, D. (1974). Connections among neurones of different types in crustacean nervous systems. In *The Neurosciences – Third Study Program*, ed. F. O. Schmitt & F. G. Worden. MIT: Cambridge, Mass.
- Koestler, A. (1967). *The Ghost in the Machine*. Hutchinson: London.
- Koestler, A. & Smythies, J. R. (1969). *Beyond Reductionism*. Hutchinson: London.
- Kortlandt, A. (1955). Aspects and prospects of the concept of instinct (vicissitudes of the hierarchy theory). *Archives néerlandaises de zoologie*, **11**, 155–284.
- Krebs, J. R. (1973). Behavioral aspects of predation. In *Perspectives in Ethology*, ed. P. P. G. Bateson & P. H. Klopfer. Plenum Press: New York & London.
- Kruuk, H. (1972). *The Spotted Hyena*. Chicago University Press: Chicago.
- Lashley, K. S. (1951). The problem of serial order in behavior. In *Cerebral Mechanisms in Behavior*, ed. L. A. Jeffres. Wiley: New York.
- Liddell, H. G. & Scott, R. (1883). *A Greek–English Lexicon*. Oxford University Press: Oxford.
- Longuet-Higgins, H. C. & Ortony, A. (1968). The adaptive memorization of sequences. In *Machine Intelligence*, vol. 3, ed D. Michie. Edinburgh University Press: Edinburgh.
- Lorenz, K. Z. (1937). Über die Bildung des Instinktbegriffes. *Naturwissenschaften*, **25**, 289–331. Reprinted (1970) as 'The establishment of the instinct concept', in *Studies in Animal and Human Behaviour*, vol. 1. Methuen: London.
- Luria, A. R. (1970). The functional organisation of the brain. *Scientific American*, **222** (3), 66–79.
- Maurus, M. & Pruscha, H. (1973). Classification of social signals in squirrel monkeys by means of cluster analysis. *Behaviour*, **47**, 106–128.
- Maynard Smith, J. & Price, G. R. (1973). The logic of animal conflict. *Nature, London*, **246**, 15–18.

- McCulloch, W. S. (1945). A heterarchy of values determined by the topology of nervous nets. *Bulletin of mathematical Biophysics*, **7**, 89–93.
- McFarland, D. J. (1971). *Feedback Mechanisms in Animal Behaviour*. Academic Press: London.
- McFarland, D. J. (1974). Time sharing as a behavioral phenomenon. In *Advances in the Study of Behavior*, **5**, 201–225.
- McFarland, D. J. & Sibly, R. M. (1975). The behavioural final common path. *Philosophical Transactions of the Royal Society of London B*, **270**, 265–293.
- Mesarovic, M. D., Macko, D. & Takahara, Y. (1970). *Theory of Hierarchical Multi-level Systems*. Academic Press: New York.
- Miller, G. A., Galanter, E. & Pribram, K. H. (1960). *Plans and the Structure of Behavior*. Holt, Rhinehart & Winston: New York.
- Myrberg, A. A. (1972). Ethology of the bicolor damselfish *Eupomacentrus partitus* (Pisces: Pomacentridae): a comparative analysis of laboratory and field behaviour. *Animal Behaviour Monographs*, **5**, 197–283.
- Nelson, K. (1964). The temporal pattern of courtship behaviour in the glandulo-caudine fishes (Ostariophysi, Characidae). *Behaviour*, **24**, 90–146.
- Nelson, K. (1973). Does the holistic study of behavior have a future? In *Perspectives in Ethology*, ed. P. P. G. Bateson & P. H. Klopfer. Plenum Press: New York & London.
- Nottebohm, F. (1970). In *Auditory Processing of Biologically Significant Sounds*, N. R. P. Bulletin, ed. F. G. Worden & R. Galambos, page 76.
- Paillard, J. (1960). The patterning of skilled movements. In *Handbook of Physiology*, ed. J. Field, H. W. Magoun & V. E. Hall, Sect. 1, vol. 3. American Physiological Society: Washington.
- Pattee, H. H. (1973). *Hierarchy Theory. The Challenge of Complex Systems*. Braziller: New York.
- Pribram, K. H. (1971). *Languages of the Brain*. Prentice-Hall: Englewood Cliffs, N.J.
- Restle, F. (1970). Theory of serial pattern learning: structural trees. *Psychological Review*, **77**, 471–495.
- Rosenblueth, A., Wiener, N. & Bigelow, J. (1943). Behavior, purpose and teleology. *Philosophy of Science*, **10**, 18–24.
- Ross, G. J. S. (1969). Algorithms AS 13, AS 14, AS 15. *Applied Statistics*, **18**, 103–106.
- Russell, E. S. (1946). *The Directiveness of Organic Activities*. Cambridge University Press: London.
- Seibt, U. (1972). Beschreibung und Zusammenspiel einzelner Verhaltensweisen von Stielaugenfliegen (Gattung *Diopsis*) unter besonderer Berücksichtigung des Putzverhaltens. *Zeitschrift für Tierpsychologie*, **31**, 225–239.
- Selfridge, O. G. & Neisser, U. (1963). Pattern recognition by machine. In *Computers and Thought*, ed. E. A. Feigenbaum & J. Feldman. McGraw-Hill: New York.
- Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, **106**, 467–482. Reprinted in Simon, H.A. (1970), *The Sciences of the Artificial*. MIT Press: Cambridge, Mass.
- Simon, H. A. (1972). Complexity and the representation of patterned sequences of symbols. *Psychological Review*, **79**, 369–382.

- Simon, H. A. (1973). The organization of complex systems. In *Hierarchy Theory*, ed. H. H. Pattee. Braziller: New York.
- Slater, P. J. B. (1973). Describing sequences of behavior. In *Perspectives in Ethology*, ed. P. P. G. Bateson & P. H. Klopfer. Plenum Press: New York & London.
- Sutherland, N. S. (1969). Outlines of a theory of visual pattern recognition in animals and man. In *Animal Discrimination Learning*, ed. R. M. Gilbert & N. S. Sutherland. Academic Press: London.
- Szentágothai, J. & Arbib, M. A. (1974). Conceptual models of neural organization. *Neurosciences Research Program Bulletin*, 12, 307-510.
- Thorpe, W. H. (1963). *Learning and Instinct in Animals*. Methuen: London.
- Thorpe, W. H. (1974). *Animal Nature and Human Nature*. Methuen: London.
- Tinbergen, N. (1950). The hierarchical organisation of nervous mechanisms underlying instinctive behaviour. *Symposia of the Society for experimental Biology*, 4, 305-312.
- Tinbergen, N. (1951). *The Study of Instinct*. Oxford University Press: London.
- Tinbergen, N. (1963). On aims and methods of ethology. *Zeitschrift für Tierpsychologie*, 20, 410-433.
- Vowles, D. M. (1970). Neuroethology, evolution and grammar. In *Development and Evolution of Behavior*, ed. L. R. Aronson, E. Tobach, D. S. Lehrman & J. S. Rosenblatt. Freeman: San Francisco.
- Weiss, P. (1941). Self-differentiation of the basic patterns of coordination. *Comparative Psychology Monographs*, 17, 1-96.
- Weiss, P. (1950). Experimental analysis of coordination by the disarrangement of central-peripheral relations. *Symposia of the Society for experimental Biology*, 4, 92-110.
- Weiss, P. (1971). *Hierarchically Organised Systems in Theory and Practice*. Hafner: New York.
- Wells, G. P. (1966). The lugworm (*Arenicola*) - a study in adaptation. *Netherlands Journal of Sea Research*, 3, 294-313.
- Winston, P. H. (1972). The M.I.T. Robot. In *Machine Intelligence*, 7, ed. B. Meltzer & D. Michie. Edinburgh University Press: Edinburgh.
- Wolff, J. G. (1975). An algorithm for the segmentation of an artificial language analogue. *British Journal of Psychology*, 66, 79-90.
- Woodger, J. H. (1937). *The Axiomatic Method of Biology*. Cambridge University Press: London.